

# Running Fluidity in Parallel

Jon Hill<sup>1</sup>

1 - Dept of Earth Science and Engineering, Imperial College London

# Outline

Pre-running

Running

Post-running

## FLML

No changes required!

[Optional]

- ▶ Remove fields from stat file
- ▶ Remove some fields from VTU

## Decompose mesh

- ▶ fldecomp
- ▶ flderecomp

Both make with `make fltools`. Programs will be in `bin`

## fldecomp

```
fldecomp -n 8 MeshName
```

Will decompose mesh into 8 peices.

## flrecomp

```
flrecomp -i 2 -n 8 InputMeshName OutputMeshName
```

Will decompose mesh from 2 to 8.

```
flrecomp -i 8 -n 2 InputMeshName OutputMeshName
```

Will decompose mesh from 8 to 2. Both need running on  
8 processors

## Local systems

```
mpiexec -n 8 ../../bin/fluidity my.flml
```

## Practical

```
cp Stommel.flml Stommel_fixed.flml
```

- ▶ Rename output
- ▶ Remove adaptivity

Run in serial and parallel



# HECToR

```
module swap PrgEnv-cray PrgEnv-fluidity
```

Submit job to back-end.

```
qsub myscript.sct
```

## HECToR

```
#!/bin/bash --login  
#PBS -N fluidity_run  
#PBS -l mppwidth=512  
#PBS -l mppnppn=4  
#PBS -l walltime=0:10:00  
#PBS -A n04-IC
```

```
module swap PrgEnv-cray PrgEnv-fluidity
```

```
# Change to the directory that the job was submitted from  
cd $PBS_O_WORKDIR
```

```
# The following take a copy of the Fluidity Python directory and  
# put it in the current directory. If we don't do this, we get import errors.
```

```
export WORKING_DIR=$(pwd -P)  
cp -r /usr/local/packages/fluidity/xe6/2.0/python/ .  
export PYTHONPATH=$WORKING_DIR/python:$PYTHONPATH
```

```
# Set the number of MPI tasks
```

```
export NPROC='qstat -f $PBS_JOBID | awk '/mppwidth/ {print $3}''
```

```
# Set the number of MPI tasks per node
```

```
export NTASK='qstat -f $PBS_JOBID | awk '/mppnppn/ {print $3}''
```

```
aprun -n $NPROC -N $NTASK fluidity -l -v1 standing_wave.flml
```

```
# clean up the python directory
```

```
rm -rf python
```

## HECToR: A run-through

- ▶ Set-up on local machine
- ▶ Copy to HECToR
- ▶ Create PBS script
- ▶ qsub
- ▶ Copy results back to local machine **or** use paraview on HECToR

## Visualisation

No different from serial - except .pvtu files, not .vtu

Log files (if used) will be one per processor.